

Leadership Computing Facility Status and Directions

NATIONAL CENTER
FOR COMPUTATIONAL SCIENCES



presented to: NCCS Users Group
March 27, 2007
Arthur S. Bland
LCF Project Director

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

Overview

- **Review of 2006**
- **LCF Funding**
- **Current Status**
- **Future**
- **Q & A**




American Competitive Initiative

In the President's State of the Union Address on January 31, 2006, President Bush stated:

“I propose to double the federal commitment to the most critical basic research programs in the physical sciences over the next ten years. This funding will support the work of America’s most creative minds as they explore promising areas such as nanotechnology, [supercomputing](#), and alternative energy sources.”

Secretary Bodman, Ibid

“Developing revolutionary, science-driven technology is at the heart of the Department of Energy’s mission. To ensure that America remains at the forefront in an increasingly competitive world, our Department is pursuing transformational new technologies in the cutting-edge scientific fields of the 21st century – areas like nanotechnology, material science, biotechnology, and [high-speed computing](#).”



Project Technical Progress Facility Upgrades

- **Site preparation for 100 TF upgrade was completed in 2nd floor computer room, on schedule**
 - Increased raised floor height to 36"
 - Installed 25 air-handlers
 - Replaced 4" pipe with 12" pipe
 - Installed 11 power distribution units




Project Technical Progress – Dual Core Upgrade

- **On-time delivery, installation, and acceptance of dual-core, 54 TF upgrade of Jaguar**
 - Swapped 5,212 processors, added 5,212 DIMMS, replaced “Y” cables in 8 days
 - 5 Optrons and 13 DIMMS had to be replaced, a failure rate of 0.2%. Remarkable for a field upgrade
 - System passed acceptance test on schedule
 - We have had substantial instability in the post-acceptance period. Cray and ORNL have worked together to get stability back close to pre-upgrade levels. (4 days between reboots vs. 6 before upgrade)




Installed 68 Cabinet XT4 System

- **First of five semi-trucks arrived on Nov. 14th - final truck arriving on Nov. 21st**
- **1,020 interconnect cables installed totaling 2.3 miles**
 - another mile of ethernet cables
 - a mile of fiber optic cables for disks!
- **No major problems with installation**
 - Minor cable problems
 - Few DIMMS replaced
- **Work completed safely on Dec. 6th**



Jaguar – 119 TF

- **Upgraded to Cray XT4**
 - Installed 68 XT4 cabinets in November
 - Moved users to XT4 system in February
 - Moved XT3 to 2nd floor and combined systems in March
- **Systems have been combined and are in acceptance testing**
- **System has 11,508 dual-core processors and 46 TB of memory**



Review of 2006 – Jaguar Operations

- **Jaguar – Cray XT3**
 - Delivered 28.4 million hours
 - Scheduled uptime of 95.7%



Review of 2006 – Phoenix Operations

- **Phoenix – Cray X1E**
 - Largest Cray vector system in the world
 - Usage of 6.3 million hours
 - System availability: 96.5%
 - Scheduled availability: 99%



Project Milestones

- **Project Start:** March 2006
- **On-time Project Completion:** September 30, 2009
- **Dates of Major Level 1 and Level 2 Milestones.** We are managing to an early-delivery schedule with schedule contingency to these dates.

	Milestone	Level 1	Level 2	Actual
✓	Acceptance of 54 TF system		9/30/2006	8/22/2006
✓	Delivery of 100 TF HW		12/31/2006	11/21/2006
✓	Acceptance of 100 TF system		3/31/2007	2/15/2007
	Delivery of 250 TF HW		12/31/2007	
	250 TF Acceptance	9/30/2008		
	Complete Site Prep for 1 PF		11/1/2008	
	Delivery of 1 PF HW		12/31/2008	
	1 PF Acceptance	9/30/2009		



Review of 2006 - Usage

CY06 INCITE and LCF Usage

Project	Jaguar			Phoenix			Type	PI
	Allocation	Usage	Remaining	Allocation	Usage	Remaining		
AST003	1,250,000	292,271	957,729	0	0	0	LCF	Burrows
AST004	3,000,000	686,929	2,313,071	0	0	0	LCF	Woosley
AST005	3,550,000	1,677,316	1,872,684	700,000	214,586	485,414	LCF	Mezzacappa
BIO014	500,000	450,966	49,034	0	0	0	LCF	Agarwal
BIO015	1,484,800	1,529,232	-44,432	0	0	0	INCITE	Karplus
CHM022	1,000,000	1,671,788	-671,788	300,000	1,004,184	-704,184	LCF	Harrison
CLI016	0	0	0	29,000	69,246	-40,246	LCF	Cessi
CLI017	3,000,000	1,995,388	1,004,612	2,000,000	1,868,508	131,492	LCF	Washington
CLI018	1,496,856	1,151,064	345,792	0	0	0	LCF	Peacock
CSC023	1,000,000	603,974	396,026	200,000	30,661	169,339	LCF	Worley
CSC026	950,000	53	949,947	0	0	0	INCITE	Smyth
EEF049	3,500,000	5,038,037	-1,538,037	300,000	553,741	-253,741	LCF	Schulthess
EEF050	0	0	0	200,000	71,319	128,681	INCITE	Hong
EEF051	500,000	153,367	346,633	0	0	0	INCITE	Ortiz
FUS011	2,000,000	6,504,969	-4,504,969	225,000	11,337	213,663	LCF	Lee
FUS012	0	0	0	440,240	880,122	-439,882	LCF	Candy
FUS013	3,000,000	1,851,943	1,148,057	0	520	-520	LCF	Batchelor
FUS014	0	0	0	400,000	872,676	-472,676	INCITE	Waltz
HEP004	30,000	0	30,000	0	0	0	LCF	Newman
HEP005	0	0	0	500,000	351,850	148,150	LCF	Ko
NPH004	1,000,000	908,521	91,479	0	0	0	LCF	Dean
SDF022	3,000,000	3,911,042	-911,042	600,000	355,502	244,498	LCF	Chen
Totals	30,261,656	28,426,860		5,894,240	6,284,252			

LCF Users and Usage

- Users come from universities, laboratories, and industry
- Usage across virtually all science domains



INCITE: 2006 and 2007



2006

- Expanded to include SC high end computing resources at PNNL, ORNL and ANL in addition to LBNL and multiple year requests.
- Received 43 proposals requesting over 95 million processor hours.
 - 60% from Universities
 - 40% had funding from other federal research agencies
- **15 awards for over 18.2 million processor hours**



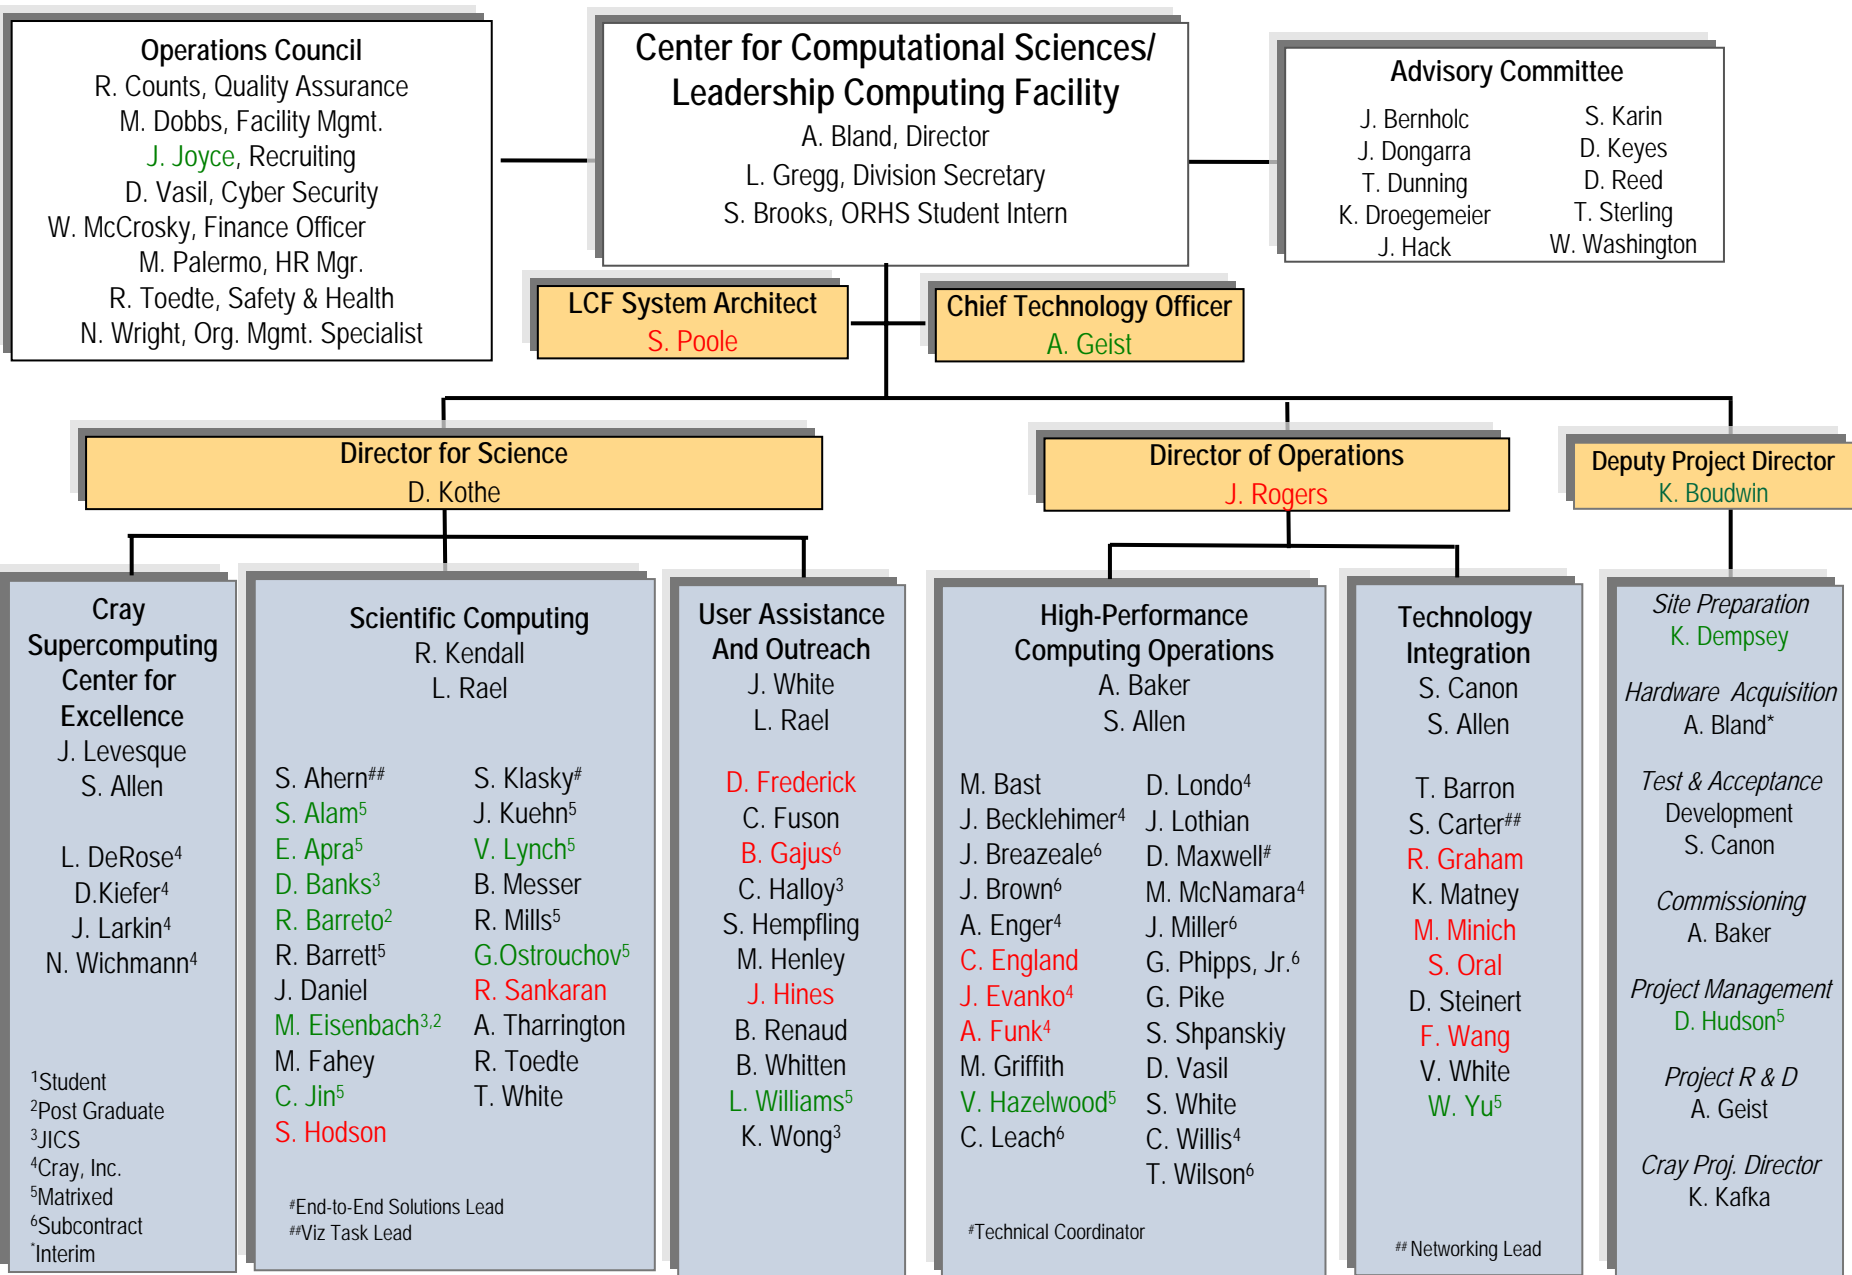
2007

- Expanded in 2007 to include 80% of resources at ORNL Leadership Computing Facilities in addition to 10% of NERSC and 5% of PNNL
- Call issued July 27, 2006
- 88 new proposals received requesting over 180 Million processor hours
- 20 renewal proposals received requesting over 75 Million processor hours
- The proposals represented the following scientific disciplines: accelerator physics, astrophysics, chemical sciences, climate research, computer science, engineering physics, environmental science, fusion energy, life sciences, materials science, nuclear physics and nuclear engineering.
- Eleven of the new proposals were from industry
- **45 INCITE awardees receive a total of over 95 Million processor hours for 2007**

2007 INCITE Allocations: 45 projects, 95 million hrs

For NCCS: 28 projects, 75 million hrs





LCF Project and Operations Budget Profile

Fiscal Year	2007		2008		2009		2010		2011	
	Cost	FTE	Cost	FTE	Cost	FTE	Cost	FTE	Cost	FTE
Approved Baseline	80.725	69	80.725	74	85.0	78	95.0	78	80.0	78
Budget Profile	77.725		77.0		85.0		95.0		80.0	

- President, DOE-SC, and Congress have been very supportive of leadership computing
- Budget realities mean that there are likely to be some hard choices in support infrastructure

Current Status - Jaguar

- **Acceptance testing has started for the combined 119 TF XT4 system**
 - Should complete testing this week and return the system to users
- **System Specs**
 - Processors: 11,508 DC
 - Memory: 2GB/core = 46 TB
 - Disk: 750 TB




Current Status - Phoenix

- **Largest Cray vector system in the world**
 - First installed in 3-2003
 - Upgraded to X1E 8-2005
- **1,024 vector processors**
- **Plan to continue operating Phoenix through 2008**
 - Based on current funding profile

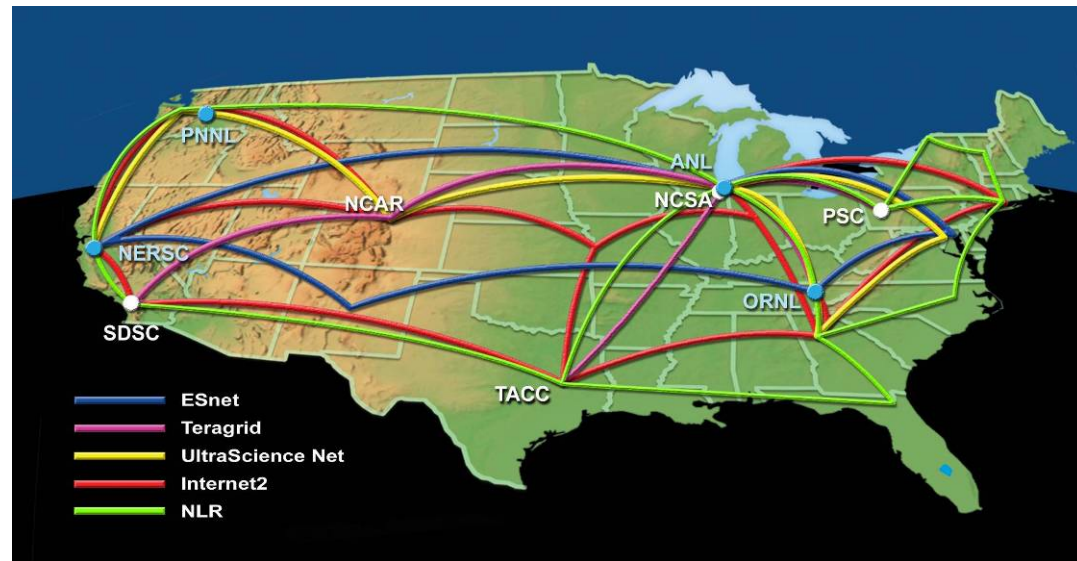


Infrastructure

- **HPSS continues strong growth**
 - Over 1.25 PB stored
 - Upgraded tape drives this year to double bandwidth
 - Adding two tape robots



- **Networks**
 - Connected to every major research network at backbone rates
 - ESnet, UltraScienceNet, Internet2, TeraGrid, GLORIAD, Cheetah




Infrastructure

- **RAM – SGI Altix**

- 256 processors
- 2 TB shared memory
- First installed in 2003
- Maintenance is on per call basis
- End of life: 9/30/2007

- **Visualization**

- “Hawk” visualization cluster
- Remote vis tools
- EVEREST powerwall



NCCS Roadmap for Leadership Computing

Mission: Deploy and operate the computational resources needed to tackle global challenges

- Future Energy
- Understanding the universe
- Nanoscale materials
- Climate Change
- Computational Biology


Vision: Maximize scientific productivity and progress on the largest scale computational problems

- Providing world class computational resources and specialized services
- Providing a stable hardware/software path of increasing scale to maximize productive applications development
- Work with users to scale applications to take advantage of systems




250 TF Upgrade

- **Jaguar will be upgraded to 250 TF in Fall 2007**
 - Replace dual-core Opteron processors with quad-core processors
 - Maintain 2 GB of memory per core
 - Each core support 4 FLOP/s per clock
- **Operating system changes to “Compute Node Linux”**
 - 6,296 4-core, 8GB shared memory nodes
 - Adds OpenMP programming model option



1000 TF System


- Cray “Baker” system is the first in Cray’s “Cascade” line designed for DARPA HPCS program
- November 2008 install
- ~ 25,000 4-core processors
- 2 GB per core
- Hypercube interconnect
- Compute Node Linux OS



Spider – Center Wide File System




**Phoenix
Cray X1E**





NFS Servers



HPSS



**ESnet, USN,
TeraGrid,
Internet2, NLR**



**Data Analysis
& Visualization**



**Jaguar
Cray XT4**



Baker


2007

- 1 PB
- 50 GB/s (aggregate)

2008

- 10 PB
- 200 GB/s (aggregate)

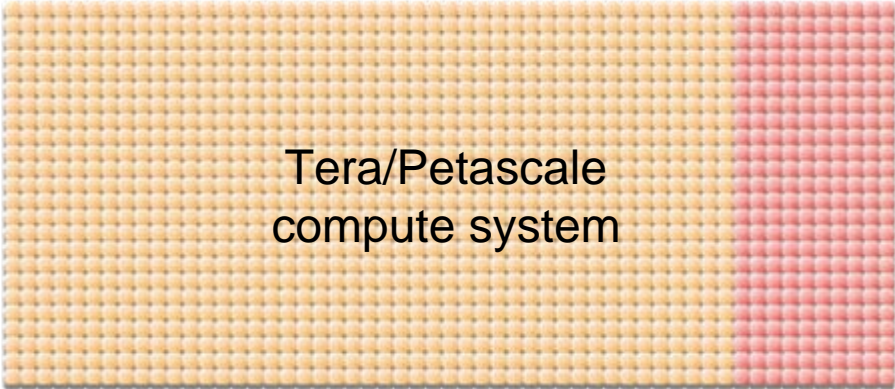
Infrastructure – Data Analysis & Visualization



“Fat nodes” support legacy visualization and analysis (IDL, AVS, etc.)



Distributed clusters support the largest analysis requirements (VisIt, EnSight, Paraview, etc.)



A large rectangular area composed of a grid of small squares. The left portion is orange and the right portion is red. The text "Tera/Petascale compute system" is centered over the orange portion.

Tera/Petascale
compute system



Moving toward a model of using a portion of the large computational resource for analysis

Our Goal is to Enable Your Science



- **Help us help you**

- Get engaged with your liaison
- We are here to help – Ask us

- **Help DOE help you**

- High expectations for breakthrough science
- We need your quarterly highlights and publications
- We have writers to turn your work into the publicity that DOE needs to keep this program going

Questions?

